
Ubiquitous Event Mining to Enhance Personal Health

Vaibhav Pandey

University of California, Irvine
Irvine, California,
United States of America
vaibhap1@uci.edu

Nitish Nag

University of California, Irvine
Irvine, California,
United States of America
nagn@uci.edu

Ramesh Jain

University of California, Irvine
Irvine, California,
United States of America
jain@ics.uci.edu

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.
UbiComp/ISWC'18 Adjunct, October 8–12, 2018, Singapore, Singapore.
Copyright © 2018 ACM ISBN/14/04...\$15.00.
<https://doi.org/10.1145/3267305.3267684>

Abstract

Advances in user interfaces, pattern recognition, and ubiquitous computing continue to pave the way for better navigation towards our health goals. Quantitative methods which can guide us towards our personal health goals will help us optimize our daily life actions, and environmental exposures. Ubiquitous computing is essential for monitoring these factors and actuating timely interventions in all relevant circumstances. We need to combine the events recognized by different ubiquitous systems and derive actionable causal relationships from an event ledger. Understanding of user habits and health should be teleported between applications rather than these systems working in silos, allowing systems to find the optimal guidance medium for required interventions. We propose a method through which applications and devices can enhance the user experience by leveraging event relationships, leading the way to more relevant, useful, and, most importantly, pleasurable health guidance experience.

Author Keywords

Event Mining, Pattern Recognition, Personal Health Navigation, User Interface

ACM Classification Keywords

C.5.3 [Computer Systems Organization]: Computer System Implementation—*Portable devices*; I.2.1 [Computer

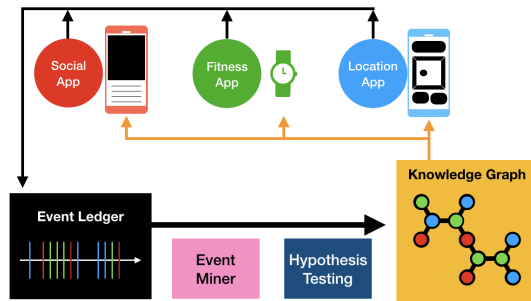


Figure 1: As devices and applications interact with a user, a timeline is created of all events pertinent to the user. Event mining establishes relationships on a temporal scale between events. These correlations provide the potential hypothesis that we may want to test on a user to produce a given outcome. After we are able to implement the hypothesis testing, we can start building a knowledge graph of events. This common set of personalized event knowledge can then be consumed by other applications or devices in real-time to continue to serve the user.

Applications]: Life and Medical Sciences; H.5.2 [Information Systems]: Information Interfaces and presentation—*User-centered design*; I.2.1 [Computing Methodologies]: Artificial Intelligence—*Medicine and Science*

Introduction

Humans have always excelled at recognizing patterns in their surroundings and utilizing them in future decision making. These patterns are observed and enriched throughout the life of an individual and shared across the community, validating some patterns as significant (or causal) and rejecting some as coincidences. This collective knowledge pool is shared and further enriched across successive generations, allowing humankind to make progress in science and technology [1].

Similar approaches can be used for ubiquitous computing and designing effective user interventions through smartphones, wearables, and IoT devices. We are constantly monitoring our life using a variety of devices and applications allowing us to estimate the state of different aspects of our life, health, and habits [4, 5]. We can take further advantage of this ubiquitous computing paradigm to develop a navigational approach to health, in which our computing services cater to our individual preferences and help actuate actions that lead us towards our health goals.

Although ubiquitous systems have operated independently and effectively in their own silos, there is still a lack of communication and knowledge sharing between these systems. By pooling the events recognized by these systems, we can discover how different aspects of our life are interconnected and how they are influencing our health state. Mining relevant patterns from the event pool can find causal relationships between these events, physiological parameters, and behaviors. Such a framework would also allow us to derive rich insights about the individual and ultimately lead to effective lifestyle interventions across various mediums of computing intervention (fig. 1). Since health is a product of all of our daily life actions, the more our computational assistance can be personalized, the better our guidance towards our health goals. One method through which we can personalize this is by ensuring that all of our devices and applications understand our historical patterns as individuals. In order for this to be useful, we need event mining to elucidate these patterns.

Finding the associations between events is the first step towards building the personal knowledge graph that can be used by different applications. We will need to implement a "do-operator" to attribute the outcome to the associated cause [6]. Once the causal relationship has been estab-

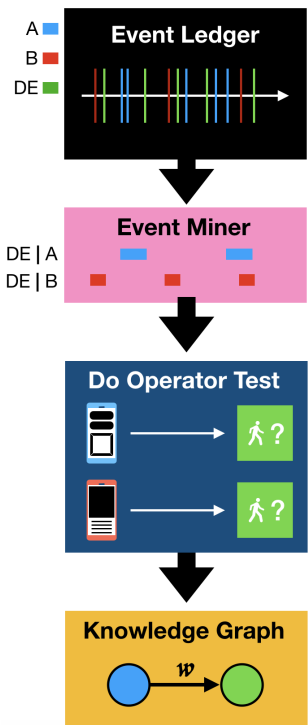


Figure 2: When we see relationships between Event A (A) or Event B (B) and a Desired Event (DE), we can test if there is a causal relationship by implementing an A/B Testing for the user. If the user responds positively towards a certain test, we can add a confidence weighted edge between two events in the knowledge graph.

lished, the event relationship can be added to the knowledge graph. Building the knowledge graph can be viewed as an exploratory step, in which we are exploring all the possible causes for an event. Similarly, guidance can be viewed as exploiting these relationships, in which different devices and applications can find a causal rule relevant in the given context and provide recommendations.

Explore: Building the knowledge graph

Every device and application should identify the associated events and write them to a shared event ledger. For example, a life logging application would find out when you are having a meal and write it to the ledger along with the associated metadata, including location, time stamp, and relevant experiential information [8], such as the name of the dish and an image if available. Figure 2 illustrates how we find the causal relationships from the event ledger and add them to the knowledge graph.

Event Miner

We can use an event mining framework as described in [2] to find associations between events. We have an event algebra to specify the relationship between different events, and we can use data driven operators to find the significant event associations from the event ledger. These event associations form the hypotheses for the do-operator. An example hypothesis is $Bike \omega_{[2,4]} Work$, which means that a bike event is followed by a work event within 2 to 4 hours.

DO-operator

The do-operator [6] takes us from observing associations to interventions. It can be most easily understood by considering the example of A/B testing in software systems. We present two different versions of the same web page to users with just one difference, such as the position of the submit button to find if the position has any effect on the

final outcome.

Similarly, in an observational dataset such as the event ledger we need to create fair datasets while controlling for the confounding factors [3] and test for the significance of the hypothesis across different datasets.

Exploit: Personal Health navigation

Richard Thaler has described in his Nobel prize winning research how our short term impulses work against the long term goals in economic decision making[7]. The same ideology could be applied to health care. Large parts of our daily habits and preferences are driven by our short term satisfaction and conveniences, instead of our long term health goals. We can find these relationships between events using the do-operators and find out how different influencing factors affect our immediate behavior. Different applications can use these relationships to "nudge"[7] the user towards their personal health goals.

Ideal Guidance Mechanism

Every individual has some preferred way of consuming information about different topics that changes how they are influenced by the information. A health navigator application can identify such behavioral preferences from the knowledge graph if we have a system monitoring the relevant aspects of the user behavior. For example, we can run an experiment using the do-operator to determine if the user's food choices are being influenced by the pictures they view and like. The navigator can use insights like these to provide the most effective guidance, which, in this instance, could be an instagram page promoting healthy food as opposed to a food blog.

Switching Between Mediums

Any application approved by the user can access the knowledge graph. The knowledge graph can also be used by

different devices to provide continuous adaptive navigation towards an individual's health goals, such as when changing mediums (leaving the car system, to a mobile phone, to a wearable). For instance, if an individual skips their morning run, then parts of their daily commute could be replaced with walking to provide the required exercise goal they had set. The knowledge graph should initially be encoded with rules derived from expert knowledge. These can be used to provide a cold-start generic recommendation while we are building the user's preference model.

Challenges

There are many research and implementation challenges in building a framework as described. We have built an event miner framework capable of finding common patterns in the data. The next steps are to design and implement the Do-operator to create fair datasets[3] with respect to the control variables and to find causal relationships between the events. We also need an event taxonomy that can be used by different devices and applications. There are also significant technical and privacy issues to consider while designing the shared event ledger.

REFERENCES

1. Yuval N. Harari. *Sapiens : a brief history of humankind*. 443 pages.
2. Laleh Jalali. 2016. Interactive Event-driven Knowledge Discovery from Data Streams. (2016).
3. Jiuyong Li, Duy Le, Lin Liu, Jixue Liu, Zhou Jin, Bingyu Sun, Saisai Ma, J Li, T D Le, L Liu, J Liu, and S Ma. 2015. From Observational Studies to Causal Rule Mining. *ACM Trans. Intell. Syst. Technol. ACM Transactions on Intelligent Systems and Technology* 7, 14 (2015), 114–123. DOI : <http://dx.doi.org/10.1145/2746410>
4. Nitish Nag, Vaibhav Pandey, Hari Bhimaraju, Srikanth Krishnan, Preston J Putzel, and Ramesh Jain. 2018. Cross-Modal Health State Estimation. *ACM Multimedia* (2018).
5. Hyungik Oh and Ramesh Jain. 2017. From Multimedia Logs to Personal Chronicles. In *Proceedings of the 2017 ACM on Multimedia Conference - MM '17*. ACM Press, New York, New York, USA, 881–889. DOI : <http://dx.doi.org/10.1145/3123266.3123375>
6. Judea Pearl and Dana Mackenzie. 2018. *The Book of Why - The New Science of Cause and Effect*. 418 pages.
7. Richard H. Thaler and Cass Sunstein. 2009. *Nudge: Improving Decisions About Health, Wealth, and Happiness*. Penguin, New York.
8. Lexing Xie, Hari Sundaram, and Murray Campbell. 2008. Event mining in multimedia streams. *Proc. IEEE* 96, 4 (2008), 623–647. DOI : <http://dx.doi.org/10.1109/JPRDC.2008.916362>